

Bayesian Group Belief

Franz Dietrich¹

London School of Economics & Maastricht University

May 2008, revised March 2010

Abstract. If a group is modelled as a single Bayesian agent, what should its beliefs be? I propose an axiomatic model that connects group beliefs to beliefs of the group members. The group members may have different information, different prior beliefs, and even different domains (algebras) within which they hold beliefs, accounting for differences in awareness and conceptualisation. As is shown, group beliefs can incorporate all information spread across individuals without individuals having to explicitly communicate their information (that may be too complex or personal to describe, or not describable in principle in the language). The group beliefs derived here take a simple multiplicative form if people's information is independent (and a more complex form if information overlaps arbitrarily). This form contrasts with familiar linear or geometric opinion pooling and the (Pareto) requirement of respecting unanimous beliefs. *JEL classification:* D70, D71

Keywords: Opinion pooling, Bayesianism, axiomatic approach, subjective probability

1 Introduction

Suppose a group is interested in whether a given hypothesis H is true. If every individual assigns a probability of 70% to H , what probability should the group as a whole assign to H ? Is it exactly 70%, or perhaps more since different persons have independently confirmed H ? The answer, I will show, crucially depends on the informational states of the individuals. If they have *identical* information, the collective has good reasons to adopt people's unanimous 70% belief, following the popular (probabilistic) Pareto principle (e.g. Mongin (1995, 1998)). Under informational asymmetry, by contrast, a possibly much higher or lower collective probability may be appropriate, and the Pareto principle becomes problematic, or so I argue.

The above question is an instance of the classic *opinion pooling/aggregation* problem, with applications for instance in expert panels. In general, the beliefs of different

¹I am very grateful for numerous helpful suggestions by a competent and diligent referee. This paper is based on my old unpublished paper 'Opinion Pooling under Asymmetric Information,' *Public Economics* 0407002, *EconWPA*, 2004. Meanwhile, interesting related results have been obtained independently by Marcus Pivato in his working paper 'The Discursive Dilemma and Probabilistic Judgement Aggregation,' *MPRA Paper* 8412, University Library of Munich, Germany, 2008.

individuals need of course not coincide, and also more than one hypothesis may be under consideration. The general goal is to merge many individuals' probability assignments to certain (exclusive and exhaustive) hypotheses into a single collective probability assignment to these hypotheses. The literature has proposed different normative conditions on the aggregation rule, and has derived the class of rules satisfying these conditions. The two most prominent types of rules are *linear* and *geometric* rules. Denoting by π_1, \dots, π_n and π the individual and collective probability assignments (each assignment being a function that maps hypotheses to probabilities), a linear rule defines π as being a weighted *arithmetic* average $\sum_{i=1}^n w_i \pi_i$, and a geometric rule defines π as being proportional to a weighted *geometric* average $\prod_{i=1}^n \pi_i^{w_i}$, where $w_1, \dots, w_n \in [0, 1]$ are fixed weights with sum 1. By contrast, our Bayesian axioms will lead to what I call *multiplicative* rules, which define π as $g \prod_{i=1}^n \pi_i$, the product of all (unweighted) individual function π_i with some fixed function g . Linear rules have been characterised (under additional technical assumptions) by the *independence* or *setwise function property* (McConway (1981), Wagner (1982, 1985), Dietrich and List (2007); see also Lehrer and Wagner (1981)), the *marginalisation property* (McConway (1981)), and (in a single-profile context) by the probabilistic analogue of the *Pareto principle* (Mongin, (1995, 1998)); and geometric rules famously satisfy *external Bayesianity* as defined in Section 6 (e.g. McConway (1978), Genest (1984), Genest, McConway and Schervish (1986)). Still an excellent reference for fundamental results on opinion pooling is Genest and Zidek's (1986) literature review.

I claim that the classic approach is problematic if, as in this paper, the goal of opinion pooling is taken to be *information aggregation*, i.e. if collective beliefs should incorporate all the information spread asymmetrically over the individuals. The classic approach is more suitable if the goal is not information aggregation: the goal might be not epistemic at all (e.g. fair representation), or it might be epistemic yet with the disagreements between individuals caused not by differences in information but by differences in interpretation of the same shared body of information.

One might at first suspect that classic pooling functions can account for informational asymmetries by putting more weight on the beliefs of well-informed individuals. More concretely, it is often suggested that in a linear and geometric rule (as defined above) the weights w_i of well-informed individuals should be higher. However, as Genest and Zidek (1986) put it, "expert weights do allow for some discrimination [...], but in vague, somewhat ill defined ways" (p. 120), and "no definite indications can be given concerning the choice or interpretation of the weights" (p. 118).

To concretely illustrate the difficulty that classic pooling functions have in aggregating information, consider again the introductory example. Suppose each individual i 's subjective probability $\pi_i(H) = 0.7$ is in fact the result of Bayesian conditioning on some private information. What should the collective belief $\pi(H)$ be? If the individuals started from the same *prior* probability of H , all depends on how this

prior compares to 0.7: if the prior is below 0.7, say 0.5, then $\pi(H)$ should intuitively exceed 0.7 because $\pi(H)$ should incorporate the pooled information of many individuals, where a single individual's information already suffices to push the probability of H up from a prior of 0.5 to a posterior of 0.7. By a similar argument, if H has a common prior above 0.7 then intuitively $\pi(H) < 0.7$, and if H has a common prior of exactly 0.7 then intuitively $\pi(H) = 0.7$. If people hold *different* prior beliefs of H , some below 0.7 and some above 0.7, then some individuals must have observed information in favour of H and the others information against H ; so, intuitively, $\pi(H)$ should be higher than 0.7 if 'most' individuals had priors of H below 0.7 (hence, had information supporting H).

These considerations highlight that knowing just the individuals' current (i.e. posterior) opinions π_1, \dots, π_n does not suffice to determine a collective opinion π that efficiently aggregates private information. But π_1, \dots, π_n are all that classic opinion pooling takes into account in calculating π . This suggests that one should depart from the classic framework. As the above example lets one suspect, the collective opinion π should be sensitive not just to people's posterior opinions π_1, \dots, π_n but also their *prior* opinions.

This paper (which is based on my unpublished paper Dietrich 2004) presents an axiomatic framework that explicitly models the information states of the individuals. The axioms lead (in the common prior case) to a unique formula for the collective probability function; no weights or other parameters are needed to incorporate all individual information into the collective beliefs. For the reason explained above, the collective beliefs depend not just on people's actual (i.e. posterior) beliefs but also their prior beliefs. This increased individual input is necessary and sufficient to efficiently aggregate information, which might come as a surprise. In short, knowing the (complex) content of people's private information is not needed: knowing people's prior-posterior pairs suffices.

As an alternative to our approach, the *supra-Bayesian* approach might also be able to aggregate information efficiently; however, despite conceptual elegance, the approach suffers from some problems, among which practical infeasibility.²

In modelling both individuals and the collective as Bayesian rationals, our findings are also relevant to the theory of *Bayesian aggregation*, which aims to merge individual beliefs/values/preferences satisfying Bayesian rationality conditions (in the sense of Savage (1954) or Jeffrey (1983)) into equally rational collective ones; for the *ex ante* approach, e.g. Seidenfeld et al. (1989), Broome (1990), Schervish et al. (1991) and

²In the supra-Bayesian approach (introduced by Morris' (1974) seminal work and extended in a large literature), collective beliefs are obtained as *posterior* probabilities (held by the real or virtual 'supra-Bayesian') conditional on the observed individual beliefs (treated as random events or evidence). This presupposes knowing (i) prior probabilities, and (ii) the likelihoods with which the individuals make probability assignments. It is not clear where these prior probabilities and likelihoods can come from; reaching a compromise or consensus on them might involve a more complex opinion pooling problem than the original one.

Mongin (1995, 1998); for the *ex post* approach, e.g. Hylland and Zeckhauser (1979), Levi (1990), Hild (1998) and Risse (2001); for an excellent overview, see Risse (2003).

Section 2 presents the axiomatic model and derives the resulting aggregation rule. Section 3 gives a numerical example. Section 4 identifies our pooling formula as a form of multiplicative opinion pooling. Sections 5 and 6 address the case of no common prior. Section 7 analyses the independent-information assumption made so far. Section 8 generalises the aggregation rule to arbitrary information overlaps.

2 An axiomatic model

Consider a group of persons $i = 1, \dots, n$ ($n \geq 2$) who need collective beliefs on certain *hypotheses*, represented as subsets H of a non-empty set Ω of *possible worlds*, i.e. worlds that are possible under the shared information. Throughout I call information (knowledge, an observation etc.) ‘shared’ if it is held by all group members. Let \mathcal{H} be the set of hypotheses $H \subseteq \Omega$ of interest, where \mathcal{H} forms a finite or countably infinite partition of Ω and $\emptyset \notin \mathcal{H}$. So, the hypotheses are mutually exclusive and exhaustive. A simple but frequent case is a binary problem $\mathcal{H} = \{H, \Omega \setminus H\}$, where H might be the hypothesis that the defendant in a court trial is guilty. In a non-binary case, \mathcal{H} might contain different hypotheses on the defendant’s extent of guilt.

In practice, the hypotheses on which opinions are formed need not be represented as subsets of a set of worlds Ω . This representation and indeed the set Ω are needed only in the present formal framework, so that we can formulate axioms, and introduce further background objects (events and probability measures) which are needed in the axioms but do not appear when applying the resulting pooling formulas.

I call an *opinion* (on \mathcal{H}) any function $f : \mathcal{H} \rightarrow (0, 1]$ with $\sum_{H \in \mathcal{H}} f(H) = 1$ (whereas *probability measures* are, as usual, defined on σ -algebras of events³); let Π be the set of all these functions f .

Let each individual i hold an opinion $\pi_i \in \Pi$, and let the collective also hold an opinion $\pi \in \Pi$. So far, this is entirely classical. Classical opinion pooling would proceed by placing conditions on how π depends on π_1, \dots, π_n , resulting in a unique relationship (e.g. $\pi = \frac{1}{n}\pi_1 + \dots + \frac{1}{n}\pi_n$) or a class of possible relationships (e.g. all linear relationships).

2.1 Simple case: common prior beliefs and a common belief domain

Before stating the axiomatic approach in full generality (that is, before allowing individuals to hold different prior beliefs defined within different domains of events),

³Any opinion uniquely extends to a probability measure defined on the σ -algebra $\sigma(\mathcal{H})$ generated by \mathcal{H} , and so we lose nothing by pooling opinions defined on \mathcal{H} rather than probability measures defined on $\sigma(\mathcal{H})$. By definition, opinions never assign zero probability to any hypothesis; this is mainly for technical convenience.

I sketch the approach in the simple base-line case. Suppose for the moment that any individual i 's opinion $\pi_i : \mathcal{H} \rightarrow (0, 1]$ is given by

$$\pi_i(H) = P(H|E_i) \text{ for all } H \in \mathcal{H},$$

where for now P is a common prior probability measure (defined on an appropriate σ -algebra over Ω , for instance the power set $\mathcal{P}(\Omega)$), and where $E_i \subseteq \Omega$ is individual i 's private information with $P(E_i) > 0$. Suppose further that people hold *independent* information: E_1, \dots, E_n are independent conditional on any hypothesis $H \in \mathcal{H}$.⁴ We would like to calculate a group opinion π . This group opinion should include all information spread over the individuals, i.e.,

$$\pi(H) = P(H|E_1 \cap \dots \cap E_n) \text{ for all } H \in \mathcal{H} \quad (1)$$

(where one easily checks that (1) is well-defined, i.e., that $P(E_1 \cap \dots \cap E_n) > 0$). One approach would be to ask all individuals i to ‘tell’ their private experience E_i , so that the group could simply gather all experiences and calculate the conditional probabilities (1). But this procedure may be unrealistic, as personal experience may be very complex and hard-to-communicate in normal language and limited time. (Another problem, which we currently assume away by using a common belief domain, is that person i 's experience E_i may be an event of which the other persons have no prior beliefs, or even no awareness or conceptualisation; asymmetries in awareness or conceptualisation might indeed explain *why* different people make different experiences.)

Assuming that private evidence cannot (or is not) communicated, can the beliefs in (1) be calculated at all? The following derivation gives a positive answer. Consider a hypothesis $H \in \mathcal{H}$ and the belief $\pi(H)$ as defined by (1). Applying Bayes' rule and then our independence assumption,

$$\begin{aligned} \pi(H) &= \frac{P(H)P(E_1 \cap \dots \cap E_n|H)}{\sum_{H' \in \mathcal{H}} P(H')P(E_1 \cap \dots \cap E_n|H')} \\ &= \frac{P(H)P(E_1|H) \cdots P(E_n|H)}{\sum_{H' \in \mathcal{H}} P(H')P(E_1|H') \cdots P(E_n|H')}. \end{aligned}$$

⁴Why do I assume that information is independent *conditional* on any hypothesis rather than unconditionally? Unconditional independence would be implausible. Suppose for instance that the information of individuals 1 and 2 both strongly correlate with the same hypothesis H in \mathcal{H} . (In a jury trial, the jurors 1 and 2 might each observe patterns in the defendant's behaviour which strongly point towards the hypothesis of guilt.) Then E_1 and E_2 are usually not independent but positively correlated ($P(E_2|E_1) > P(E_2)$), because learning E_1 raises the probability of H , which in turn raises that of E_2 . More generally, since the evidences E_1, \dots, E_n tell something about the hypotheses, learning *some* of the E_i s leads to revised probabilities of the hypotheses, which leads to revised probabilities of the other E_i s. In short, the E_i s are non-independent because they are mutually relevant *via* their relevance to hypotheses in \mathcal{H} . This argument for non-independence is blocked once we condition on a hypothesis: conditional on a given hypothesis being true, evidences are not relevant to (i.e., do not bring new information about) hypotheses. If all existing probabilistic dependence between evidences goes ‘via’ the hypotheses, then conditioning on a hypothesis eliminates all sources of dependence, and the evidences become conditionally independent. Our (conditional) independence assumption is analysed again below.

In the numerator and the denominator, each factor of type $P(E_i|H)$ can be rewritten according to

$$P(E_i|H) = \frac{P(H|E_i)P(E_i)}{P(H)} = \frac{\pi_i(H)P(E_i)}{P(H)}.$$

Substituting this expression, we obtain

$$\begin{aligned} \pi(H) &= \frac{P(H) \frac{\pi_1(H)P(E_1)}{P(H)} \cdots \frac{\pi_n(H)P(E_n)}{P(H)}}{\sum_{H' \in \mathcal{H}} P(H') \frac{\pi_1(H')P(E_1)}{P(H')} \cdots \frac{\pi_n(H')P(E_n)}{P(H')}} \\ &= \frac{\pi_1(H) \cdots \pi_n(H) / P(H)^{n-1}}{\sum_{H' \in \mathcal{H}} \pi_1(H') \cdots \pi_n(H') / P(H')^{n-1}}. \end{aligned}$$

Interestingly, any private information E_i has dropped out altogether, so that the collective opinion π can be calculated solely on the basis of the revealed individual opinions π_1, \dots, π_n (and the fixed prior). Put differently, each individual information E_i has been incorporated without disclosing it. In short, denoting by p the prior opinion $P|_{\mathcal{H}}$ (i.e., the restriction of P to the hypotheses of interest), we have shown that

$$\pi \propto \pi_1 \cdots \pi_n / p^{n-1}.$$

Here and throughout, I call functions $f, g : \mathcal{H} \rightarrow \mathbb{R}$ *proportional*, written $f \propto g$, if there exists a constant $k \neq 0$ such that $f(H) = kg(H)$ for all $H \in \mathcal{H}$.

2.2 General case: possibly distinct prior beliefs and belief domains

After this preliminary analysis, let us start afresh, this time in full generality, and stating all assumptions as explicit axioms. Recall that we consider individual opinions $\pi_1, \dots, \pi_n \in \Pi$ and a collective opinion $\pi \in \Pi$. The further elements introduced in the preliminary Section 2.1 (namely, P, E_1, \dots, E_n) are now re-introduced in their general and official form. For each person i there is (without having to be revealed):

- an event $E_i \subseteq \Omega$, i 's *personal information*;
- a ('prior') probability measure P_i representing i 's beliefs based on the *shared* information (hence prior to observing E_i). P_i need not assign a probability to *all* events in $\mathcal{P}(\Omega)$; rather, P_i is defined on some σ -algebra $\mathcal{A}_i \subseteq \mathcal{P}(\Omega)$, containing the events on which i holds beliefs (whereas on other events i may lack beliefs, or even lack awareness or conceptualisation). But \mathcal{A}_i should contain at least E_i and all hypotheses in \mathcal{H} , where $P_i(E_i) > 0$ and $P_i(H) > 0$ for all $H \in \mathcal{H}$. The restriction of i 's prior belief P_i to \mathcal{H} is called i 's *prior opinion*. It is denoted by $p_i \in \Pi$ and given by $p_i(H) = P_i(H)$ for all $H \in \mathcal{H}$.

These model resources allow us to state a standard rationality condition:

Individual Bayesian Rationality (IBR) $\pi_i(H) = P_i(H|E_i)$ for each person i and hypothesis $H \in \mathcal{H}$.⁵

A person i 's belief domain \mathcal{A}_i may fail to contain another person j 's observation E_j , and this for (at least) two reasons. First, the fact that j but not i observed E_j may be due precisely to j having subjectively conceptualised E_j but i not having done so; juror j in a trial may be the only juror to observe the suspicious smile on the defendant's face *because* the other jurors i do not even know what a suspicious smile would be like. Second, j 's information E_j may be so detailed and complex that prior to j observing it, it belonged not even to j 's own belief domain, let alone to i 's; that is, it was only while observing E_j that person j extended his prior beliefs to a larger domain \mathcal{A}_j containing E_j .

Following the paradigm of social choice theory, I treat the collective as a separate virtual agent with its own beliefs. While this agent is typically a construction (i.e. there needn't exist any real individual holding these beliefs), the social choice paradigm requires it to be as rational as any real individual.⁶ 'Rationality' refers to different things in different contexts (e.g. to transitivity of preferences in Arrovian preference aggregation, to von-Neumann-Morgenstern rationality in Harsanyi's Theorem on group preferences over lotteries, to logical consistency in judgment aggregation, and so on). In the present context, it naturally refers to Bayesian rationality. To formulate this, I suppose that there is

- a ('prior') probability measure P , representing collective beliefs based on people's *shared* information (hence not on their private information E_1, \dots, E_n). P is defined on some σ -algebra $\mathcal{A} \subseteq \mathcal{P}(\Omega)$, the domain of the collective beliefs, which contains at least all private evidences E_1, \dots, E_n and all $H \in \mathcal{H}$, where $P(E_1 \cap \dots \cap E_n) > 0$ and $P(H) > 0$ for all $H \in \mathcal{H}$. The restriction of the collective prior belief P to \mathcal{H} is called the *collective prior opinion*; it is denoted by p ($\in \Pi$) and given by $p(H) = P(H)$ for all $H \in \mathcal{H}$.

\mathcal{A} , P and p are the collective counterparts of \mathcal{A}_i , P_i and p_i . The collective counterpart of (IBR) is:

Collective Bayesian Rationality (CBR) $\pi(H) = P(H|E_1 \cap \dots \cap E_n)$ for each hypothesis $H \in \mathcal{H}$.

⁵The conditional probability $P_i(H|E_i)$ is well-defined because $E_i, H \in \mathcal{A}_i$ and $P_i(E_i) > 0$. Our assumptions also take care that all other conditional probabilities used in this paper are well-defined.

⁶The collective agent should be rational notably because it forms the basis for collective actions and decisions.

Condition (CBR) requires the collective opinion π to incorporate all information spread over people: the *shared* information (contained in the prior P) and all *personal* information (contained in E_1, \dots, E_n).

While we have ensured, via (CBR), that the collective opinion uses all evidence scattered across individuals, we have done nothing so far to constrain the collective prior probability measure P (which underlies π). Indeed, P may so far be totally disconnected from the individual prior probability measures P_1, \dots, P_n (which underlie π_1, \dots, π_n). The next condition does something to connect P to P_1, \dots, P_n . More precisely, the condition ties the likelihood that the collective assigns to the various individual evidences E_1, \dots, E_n to the individuals' own likelihood assessments:

Accept People's Likelihood Assessments (APLA) For all persons i and hypotheses $H \in \mathcal{H}$, $P(E_i|H) = P_i(E_i|H)$.

This principle requires the collective to take over i 's own interpretation of i 's information E_i as given by i 's likelihood assignments $P_i(E_i|H)$, $H \in \mathcal{H}$. To motivate this condition, let me first explain the context in a little more detail. In statistics, the information that data contain on given hypotheses (as opposed to prior beliefs on these hypotheses) is usually taken to be summarised in the data's likelihood function, which maps any hypothesis to the data's probability given this hypothesis. For instance, the information on humidity contained in a temperature measurement of 20 degrees Celsius is given by the mapping that assigns to each potential humidity level the probability that temperature is 20 degrees Celsius given this humidity level. In our case, the information contained in individual i 's evidence E_i is summarised in E_i 's likelihood function, mapping any hypothesis H to E_i 's probability given H . But how large exactly is E_i 's probability given H ? For instance, how probable is it that the defendant in a trial has a particular facial expression (E_i) given the hypothesis that he is guilty (H)? The answer may be far from trivial, as one might come up with various different interpretations of the same observation. Condition (APLA) requires that the answer that the collective gives matches the answer that the individual *who observed the evidence* gives; that is, $P(E_i|H) = P_i(E_i|H)$. What is the motivation behind identifying $P(E_i|H)$ with $P_i(E_i|H)$? Why not also take other persons' interpretations of E_i into account by defining $P(E_i|H)$ as some compromise of $P_1(E_i|H), \dots, P_n(E_i|H)$? First, for reasons explained above, the persons $j \neq i$ may not even hold beliefs on the unobserved event E_i (i.e., $E_i \notin \mathcal{A}_j$), in which case $P_j(E_i|H)$ is simply undefined. Second, assuming that the persons $j \neq i$ do hold such beliefs (i.e., $E_i \in \mathcal{A}_j$), a 'likelihood compromise' could only be formed after each person j reveals $P_j(E_i|H)$; which in turn supposes that first i communicates his informational basis E_i in all detail to the rest of the group. This is not only at odds with the present approach, but may also be infeasible: given the possible complexity of E_i and the limitations of language, time, i 's ability to describe E_i , j 's ($j \neq i$) ability to understand E_i , and so on, j could probably learn at most some approximation \tilde{E}_i of

E_i , and so j could at most provide j 's likelihood of \tilde{E}_i , which only approximates j 's likelihood of the true E_i ($P_j(\tilde{E}_i|H) \approx P_j(E_i|H)$).

The next assumption is not a normative condition but rather an assumption on the environment: individuals receive *independent* information. This assumption will be analysed (and relaxed) in later sections; see footnote 4 above for first considerations. For now, I only mention that it is strong but very common. It is analogous to (i) independence assumptions on private information/types in Bayesian games, (ii) the independence condition in the literature on the Condorcet Jury Theorem (see Dietrich 2008 for a critique of the condition), (iii) the *Parental Markov Condition* in the theory of Bayesian networks (interpreting the true hypothesis in \mathcal{H} as the parent of each information E_i in a Bayesian network; see Pearl 2000), and (iv) Fitelson's (2001) condition of *confirmational independence*.

Independent Information (Ind) For each hypothesis $H \in \mathcal{H}$, the personal observations E_1, \dots, E_n are independent conditional on H .⁷

I am ready to state the theorem. Recall that p_i, π_i is the pair of person i 's prior and posterior opinion, and p, π is the pair of the collective prior and posterior opinion.

Theorem 1 *Suppose individuals satisfy (IBR), information satisfies (Ind), and the collective satisfies (CBR) and (APLA). Then the collective opinion π is proportional to the collective prior opinion times all individual posterior-to-prior ratios:*

$$\pi \propto p \frac{\pi_1}{p_1} \dots \frac{\pi_n}{p_n}.$$

Proof. Suppose (IBR), (CBR), (APLA) and (Ind) hold. For all H in \mathcal{H} ,

$$\begin{aligned} \pi(H) &= P(H|E_1 \cap \dots \cap E_n) \text{ by (CBR)} \\ &= \frac{P(E_1 \cap \dots \cap E_n|H)p(H)}{P(E_1 \cap \dots \cap E_n)} \text{ by Bayes' rule} \\ &= kP(E_1 \cap \dots \cap E_n|H)p(H) \text{ for a constant } k \neq 0 \\ &= kP(E_1|H) \dots P(E_n|H)p(H) \text{ by (Ind)} \\ &= kP_1(E_1|H) \dots P_n(E_n|H)p(H) \text{ by (APLA)} \\ &= k \frac{P_1(H|E_1)P_1(E_1)}{p_1(H)} \dots \frac{P_n(H|E_n)P_n(E_n)}{p_n(H)} p(H) \text{ by Bayes' rule} \\ &= k' \frac{P_1(H|E_1)}{p_1(H)} \dots \frac{P_n(H|E_n)}{p_n(H)} p(H) \text{ for a constant } k' \neq 0 \\ &= k' \frac{\pi_1(H)}{p_1(H)} \dots \frac{\pi_n(H)}{p_n(H)} p(H) \text{ by (IBR). } \blacksquare \end{aligned}$$

Three important remarks are due.

⁷As usual, by 'independence' of events I mean full independence, not just pairwise independence.

1. As promised, the collective opinion π is calculated without people having to communicate their arbitrarily complex informational bases E_i or their likelihoods $P(E_i|H)$, $H \in \mathcal{H}$. In practice, all persons i submit their prior-posterior pairs p_i, π_i (or just their ratios π_i/p_i), and then the collective opinion π is calculated. So, compared to standard opinion pooling, we additionally require submission of prior opinions p_1, \dots, p_n , a complication that enables the incorporation of the individual information E_1, \dots, E_n into the collective opinion.

2. Theorem 1's formula does not fully solve the aggregation problem since we do not yet know how to determine the collective prior opinion p . Strategies to choose p are presented in Sections 5 and 6. In practice, there is an alternative to having to choose p : one might use an *approximation* of Theorem 1's formula, defining the collective opinion as

$$\pi_{\text{approx}} \propto \frac{\pi_1}{p_1} \cdots \frac{\pi_n}{p_n}, \quad (2)$$

the product of all posterior-to-prior ratios. When and why can π_{approx} count as a good approximation of Theorem 1's formula? Let me give some heuristic arguments.⁸ Often, the function $\frac{\pi_1}{p_1} \cdots \frac{\pi_n}{p_n}$ varies considerably, i.e., assigns very different values to the hypotheses H in \mathcal{H} . Intuitively, this is because pooled information is often strong evidence for or against certain hypotheses. More formally, if sufficiently many of the individual ratios $\frac{\pi_i}{p_i}$ vary at least moderately, the product $\frac{\pi_1}{p_1} \cdots \frac{\pi_n}{p_n}$ varies strongly (provided that the individual variations do not systematically cancel each other out).⁹ Whenever the variation of $\frac{\pi_1}{p_1} \cdots \frac{\pi_n}{p_n}$ is strong enough to 'outweigh' that of p (assuming p should not vary very much), the function $\frac{\pi_1}{p_1} \cdots \frac{\pi_n}{p_n}$ varies roughly like $p^{\frac{\pi_1}{p_1}} \cdots p^{\frac{\pi_n}{p_n}}$; and hence, the opinions π_{approx} and π (obtained by normalising the two functions so as to each sum to one) are roughly similar.

3. Assume a unanimous *posterior* agreement $\pi_1 = \dots = \pi_n$ (as in the introduction's example). Then only in special cases does π equal $\pi_1 = \dots = \pi_n$, which shows that the unanimity/Pareto principle often required in standard opinion pooling is problematic under informational asymmetries. One such special case is that $\pi_1 = \dots = \pi_n = p_1 = \dots = p_n = p$, so that none of the personal observations E_1, \dots, E_n confirms or disconfirms any hypothesis, i.e., in essence, there is no informational asymmetry.

An important special case of Theorem 1 is that where people have managed to agree on how to interpret their shared information, i.e. where they hold a common prior opinion:

Common Prior (CP) $p_1 = \dots = p_n = p$ (i.e., the prior probability measures P_1, \dots, P_n, P agree on all hypotheses in \mathcal{H} , though perhaps not elsewhere).

⁸I owe these thoughts to the helpful referee.

⁹If for instance most individual ratios peak at the same hypothesis (say, if most jurors believe the defendant is guilty) then the product of ratios is likely to strongly peak at this hypothesis.

Corollary 1 *Under the assumptions of Theorem 1 and (CP), the collective opinion π is given by*

$$\pi \propto \pi_1 \cdots \pi_n / p_1^{n-1}.$$

Let me make three remarks on this corollary.

1. The corollary’s formula differs in an important respect from Theorem 1’s formula: the parameter p has been eliminated, and so the collective opinion π is fully determined by the individual prior and posterior opinions. By contrast, if (CP) fails, i.e. if the group didn’t manage to agree on how to interpret the shared information, Theorem 1’s formula does not fully solve the aggregation problem, as we need a way to determine the collective prior p (see Sections 5 and 6).

2. Condition (CP) can in fact be seen as the conjunction of two conditions. The first (descriptive) condition is that $p_1 = \dots = p_n$, i.e. all persons i submit the same prior opinion. The second (normative) condition is that the unanimity (or Pareto) principle holds for the *prior* opinions, i.e. if all submit the same prior opinion, this becomes the collective prior opinion. Applying a unanimity condition to prior opinions is far less problematic than doing so for the posterior opinions π_1, \dots, π_n, π , because prior opinions contain no informational asymmetry.

3. According to a prominent view, held notably by Harsanyi, *any* inter-personal differences in beliefs between rational agents stem from different *information* (in a suitably general sense of this word), never from different prior beliefs.¹⁰ If this view is correct, and if each opinion p_i indeed incorporates no information except the shared one which does not depend on i , then the p_i s must be identical. So, Harsanyi’s view places us in the comfortable position of being able to assume (CP). Harsanyi’s view is based on modelling all experiences – including early ones in life and perhaps prenatal ones – as information shaping beliefs via Bayesian conditioning. If we think of i ’s private information E_i as containing all such experiences, then the events E_i (and the underlying space Ω) inevitably become highly complex. While this by itself poses no problem (since E_1, \dots, E_n need not be revealed), at least one difficulty arises. Even if Harsanyi were fundamentally right, people will in practice often not agree on a common prior opinion, if only because they do not ‘remember’ the common prior opinion that they used to hold at the prenatal stage. I personally do not share Harsanyi’s view. I believe in the possibility of genuinely non-information-driven disagreements, and hence in the possibility that p_1, \dots, p_n differ in spite of incorporating the same (shared) information.¹¹

¹⁰I am grateful to the referee for raising this issue.

¹¹Harsanyi’s claim is true nearly by definition if the notion of ‘information’ is purely technical and if the claim is taken to be one about *mathematically representing* different probability measures as being obtained by conditioning from a common probability measure (defined on a suitably extended algebra of events). Under a so-extended notion of ‘information’, the whole process of personal deliberation needed to form one’s beliefs and to interpret one’s information constitutes another large piece of (‘meta-’)information. The current paper’s notion of ‘information’ is not of this abstract kind. It

4. Instead of interpreting E_i as reflecting all of i 's personal information, one might re-interpret E_i as reflecting only that part of i 's personal information which i has incorporated rationally into his opinion (in the sense of Bayesian conditioning). Then (IBR) becomes true by definition. All not rationally incorporated personal information is then simply thrown away, i.e., not included in collective beliefs.¹²

3 A numerical example for a simple case

Consider the simple case of a binary problem $\mathcal{H} = \{H, \Omega \setminus H\}$ (H and $\Omega \setminus H$ might mean that the defendant in a court trial is guilty resp. innocent, and persons might be jurors). Suppose Common Prior (CP), i.e. $p_1 = \dots = p_n = p$. By Theorem 1 (that is, by its corollary), the collective posterior of H is given by

$$\begin{aligned} \pi^H &= \frac{\pi_1^H \dots \pi_n^H / (p^H)^{n-1}}{\pi_1^H \dots \pi_n^H / (p^H)^{n-1} + \frac{(1 - \pi_1^H) \dots (1 - \pi_n^H)}{1} / (1 - p^H)^{n-1}} \\ &= \frac{1}{1 + (1/\pi_1^H - 1) \dots (1/\pi_n^H - 1) / (1/p^H - 1)^{n-1}}, \end{aligned} \quad (3)$$

where $p^H := p(H)$, $\pi^H := \pi(H)$ and $\pi_i^H := \pi_i(H)$. For the case of only $n = 2$ individuals, in which the formula (3) for the collective posterior reduces to

$$\begin{aligned} \pi^H &= \frac{\pi_1^H \pi_2^H / p^H}{\pi_1^H \pi_2^H / p^H + \frac{(1 - \pi_1^H)(1 - \pi_2^H)}{1} / (1 - p^H)} \\ &= \frac{1}{1 + (1/\pi_1^H - 1)(1/\pi_2^H - 1) / (1/p^H - 1)}, \end{aligned}$$

Table 1 contains the values of the collective belief π^H for all possible combinations of values of p^H, π_1^H, π_2^H in the grid $\{0.1, 0.25, 0.5, 0.75, 0.9\}$.¹³ Note how drastically the group belief π^H depends on the prior p^H . By shifting p^H below (above) the individual posteriors π_i^H , the group belief π^H quickly approaches 1 (0). The interpretation is that if the posteriors π_i^H are all to the same side of the prior, then the evidences

is a substantive notion, under which it is possible that two opinions incorporate no (or *the same*) information and yet differ. If however Harsanyi's claim is taken to be not just about mathematical representability but about the psychological reality of rational agents, then the claim is problematic. Many experiences in life (such as hearing a sound for the first time) have a content that was not previously conceptualized by the agent, hence cannot belong to the algebra within which the agent previously held beliefs. So the agent's new beliefs after the experience cannot stem from updating the old beliefs by Bayesian conditioning on this event. The topic of non-informational belief formation goes beyond this paper (but will be developed in the paper 'A reason-based theory of rational belief' with Christian List).

¹²Instead of throwing this information away, one might ask people to incorporate it (in some non-Bayesian ways, unfortunately) in the submitted prior opinions. This removes the informational symmetry underlying the submitted prior opinions, which in turn affects the interpretation and plausibility of the analysis of later sections.

¹³The entries are rounded results if three decimal digits are reported, and exact results else.

		$p^H :$				
		0.1	0.25	0.5	0.75	0.9
$\pi_1^H, \pi_2^H :$	0.1, 0.1	0.1	.036	.012	.004	.001
	0.25, 0.1	0.25	0.1	.036	.012	.004
	0.25, 0.25	0.5	0.25	0.1	.036	.012
	0.5, 0.1	0.5	0.25	0.1	.036	.012
	0.5, 0.25	0.75	0.5	0.25	0.1	.036
	0.5, 0.5	0.9	0.75	0.5	0.25	0.1
	0.75, 0.1	0.75	0.5	0.25	0.1	.036
	0.75, 0.25	0.9	0.75	0.5	0.25	0.1
	0.75, 0.5	0.964	0.9	0.75	0.5	0.25
	0.75, 0.75	0.988	0.964	0.9	0.75	0.5
	0.9, 0.1	0.9	0.75	0.5	0.25	0.1
	0.9, 0.25	0.964	0.9	0.75	0.5	0.25
	0.9, 0.5	0.988	0.964	0.9	0.75	0.5
	0.9, 0.75	0.996	0.988	0.964	0.9	0.75
	0.9, 0.9	0.999	0.996	0.988	0.964	0.9

Table 1: Collective probability $\pi^H = \pi(H)$ in dependence of the common prior $p^H = p(H)$ and the individual posteriors $\pi_i^H = \pi_i(H)$, for a group of size $n = 2$.

E_i all point into the same direction, so that their conjunction points even more into that direction. But if the prior p^H is somewhere in the middle of the posteriors π_i^H , the group belief π^H may be moderate. The interpretation is that if some of the posteriors π_i^H are above the prior and others are below the prior, then the evidences E_i point into different directions, and their conjunction need not strongly point into any direction. The above formula for the group belief π^H shows that it strictly increases as a function of each individual belief π_i^H , but strictly decreases as a function of the prior belief p^H . But how can one make sense of the group posterior π^H depending negatively on the prior p^H ? How could more prior support for H possibly reduce H 's posterior probability? The answer is that by increasing the prior p^H while keeping the individual posteriors π_i^H fixed one implicitly reduces the support that each of the evidences E_i gives to H ; as a result, the collective posterior of H falls, intuitively because the reduced evidential support for H overcompensates the increased prior support.

4 Multiplicative opinion pooling

If we treat the prior opinions p_1, \dots, p_n, p as fixed parameters, the pooling formula of Theorem 1 depends just on π_1, \dots, π_n , hence defines a classic pooling function $F : \Pi^n \rightarrow \Pi$. Specifically, this pooling function is given by $\pi = g \cdot \pi_1 \cdots \pi_n$ where g is

a fixed function on \mathcal{H} given by $g \propto p/(p_1 \cdots p_n)$ (and in particular as $g \propto p^{1-n}$ under Common Prior (CP)). So, our axioms lead to what one might call a *multiplicative* opinion pool. Formally, a (classic) opinion pool $F : \Pi^n \rightarrow \Pi$ is *multiplicative* if it is given by

$$F(\pi_1, \dots, \pi_n) = g \cdot \pi_1 \cdots \pi_n \text{ for all } \pi_1, \dots, \pi_n \in \Pi,$$

for some fixed function $g : \mathcal{H} \rightarrow (0, \infty)$. The simplest multiplicative rule is that in which g is a constant function, so that

$$F(\pi_1, \dots, \pi_n) \propto \pi_1 \cdots \pi_n \text{ for all } \pi_1, \dots, \pi_n \in \Pi.$$

Note how multiplicative opinion pools differ from the more common linear and geometric opinion pools; these arise from different axiomatic systems that do not make information explicit.

In fact, our axioms not only imply that pooling is multiplicative: they *characterise* multiplicative pooling if \mathcal{H} is finite because every multiplicative rule can be obtained from suitable priors $p_1, \dots, p_n, p \in \Pi$.¹⁴

5 Choosing the collective prior p when there is no common prior

If the interpretation of the shared information is controversial and hence (CP) fails, the group needs to determine the collective prior p in Theorem 1's formula. At least three strategies are imaginable. First, one might define p as a uniform or maximum-entropy prior if available. Second, someone, not necessarily a group member, may be appointed to choose p , either by drawing on his own prior beliefs, or by taking inspiration from the submitted priors p_1, \dots, p_n , or by using statistical estimation techniques if available. These two solutions have obvious limitations, including some ad-hoc-ness and a lack of democracy. A third alternative is to replace p by $F(p_1, \dots, p_n)$ and thus define the collective opinion by

$$\pi \propto \frac{\pi_1}{p_1} \cdots \frac{\pi_n}{p_n} F(p_1, \dots, p_n), \quad (4)$$

where $F : \Pi^n \rightarrow \Pi$ is a standard opinion pool. Note that F is used here not to aggregate people's actual (posterior) opinions π_1, \dots, π_n but to aggregate their prior opinions p_1, \dots, p_n , namely into a 'compromise prior'. At first sight, one may wonder what is gained by formula (4) compared to the standard approach of defining $\pi = F(\pi_1, \dots, \pi_n)$ without having to care about priors p_1, \dots, p_n . Does formula (4) not just shift the classic aggregation problem – pooling π_1, \dots, π_n into π – towards an equally

¹⁴For any multiplicative rule $F : \Pi^n \rightarrow \Pi$, say generated by the function g , if we (for instance) take p_1, \dots, p_n, p to be all identical and proportional to $g^{-1/(n-1)}$, then $g \propto p/(p_1 \cdots p_n)$, and hence the multiplicative rule generated by g coincides with that arising in Theorem 1.

complex aggregation problem about priors – pooling p_1, \dots, p_n into p ? In an important respect, pooling p_1, \dots, p_n is simpler than pooling π_1, \dots, π_n : unlike π_1, \dots, π_n , the prior opinions p_1, \dots, p_n involve no informational asymmetry since each p_i is based on the same (shared) information.¹⁵ Hence any disagreement between p_1, \dots, p_n is due solely to different *interpretations* of that same body of information. This may facilitate the choice of F . For instance, aggregation may be guided by the unanimity/Pareto principle (which is problematic under informational asymmetry, as we have seen). Further, aggregation may place *equal weights* on each of the priors p_1, \dots, p_n (whereas pooling π_1, \dots, π_n may involve the difficult and vague exercise of assigning more weight to better informed people). The literature’s two most prominent types of opinion pools $F : \Pi^n \rightarrow \Pi$ are

$$\begin{aligned} \text{linear opinion pools:} \quad & F(p_1, \dots, p_n) = w_1 p_1 + \dots + w_n p_n, \\ \text{geometric opinion pools:} \quad & F(p_1, \dots, p_n) \propto p_1^{w_1} \dots p_n^{w_n}, \end{aligned}$$

with weights $w_1, \dots, w_n \in [0, 1]$ that add up to 1 (where in the geometric pool the factor of proportionality is chosen such that $\sum_{H \in \mathcal{H}} F(p_1, \dots, p_n)(H) = 1$). If F is a linear resp. geometric opinion pool, our pooling formula (4) becomes

$$\pi = \frac{\pi_1}{p_1} \dots \frac{\pi_n}{p_n} (w_1 p_1 + \dots + w_n p_n) \quad (5)$$

$$\text{resp. } \pi \propto \frac{\pi_1}{p_1} \dots \frac{\pi_n}{p_n} p_1^{w_1} \dots p_n^{w_n} = \frac{\pi_1}{p_1^{1-w_1}} \dots \frac{\pi_n}{p_n^{1-w_n}}. \quad (6)$$

How should the weights w_1, \dots, w_n be chosen in practice? In general, unequal weights may be justified *either* by different information states *or* by different competence, i.e. ability to interpret information. The former reason does not apply here, since p_1, \dots, p_n are by definition based on the same (shared) information. If, in addition, differences of competence are either inexistent, or unknown, or not to be taken into account for reasons of procedural fairness, then equal weights $w_1 = \dots = w_n = 1/n$ are justified, so that our pooling formula becomes

$$\pi = \frac{1}{n} \frac{\pi_1}{p_1} \dots \frac{\pi_n}{p_n} (p_1 + \dots + p_n) \quad (7)$$

$$\text{resp. } \pi \propto \frac{\pi_1}{p_1^{1-1/n}} \dots \frac{\pi_n}{p_n^{1-1/n}}, \quad (8)$$

which is parameter-free, hence uniquely solves the aggregation problem.

6 External and internal Bayesianity

I now give an argument in defence of defining F in (4) as a geometric (or more generally, externally Bayesian) opinion pool, hence in defence of our pooling formulae

¹⁵One might even argue that, while pooling p_1, \dots, p_n into p is possible without using extra information (due to the informational symmetry), pooling π_1, \dots, π_n into π is impossible without extra information (such as p_1, \dots, p_n).

(6) and (8). Note first that in (4) π is a function of the vector $(p_1, \pi_1, \dots, p_n, \pi_n) \in (\Pi \times \Pi)^n = \Pi^{2n}$, containing every person's prior and posterior.

Definition 1 A generalised opinion pool ('GOP') or generalised probability aggregation rule is a function $G : \Pi^{2n} \rightarrow \Pi$.

Unlike a standard opinion pool $F : \Pi^n \rightarrow \Pi$, a GOP G also takes as inputs the p_i s, i.e. people's interpretations of the shared information. As shown above, our axioms imply that a GOP G should take the form (4), i.e. the form

$$G(p_1, \pi_1, \dots, p_n, \pi_n) \propto \frac{\pi_1}{p_1} \cdots \frac{\pi_n}{p_n} F(p_1, \dots, p_n) \quad (9)$$

where $F : \Pi^n \rightarrow \Pi$ is a standard opinion pool that merges the priors p_1, \dots, p_n .

From a Bayesian perspective, two natural conditions may be imposed on a GOP, to be called *external* and *internal Bayesianity*. The former is an analogue of the equally-named classic condition for standard opinion pools F : it should not matter whether information arrives before or after pooling, i.e. pooling should commute with Bayesian updating. Formally, for every opinion $p \in \Pi$ and (likelihood) function $l : \mathcal{H} \rightarrow (0, 1]$ the (updated) opinion $p^l \in \Pi$ is defined by

$$p^l(H) := \frac{l(H)p(H)}{\sum_{H' \in \mathcal{H}} l(H')p(H')}, \text{ in short } p^l \propto lp. \quad (10)$$

Here, l is interpreted as a likelihood function $P(E|\cdot)$ for some observation E , so that p^l is a posterior probability. A standard opinion pool $F : \Pi^n \rightarrow \Pi$ is called *externally Bayesian* if

$$F(p_1^l, \dots, p_n^l) = F(p_1, \dots, p_n)^l$$

for every profile $(p_1, \dots, p_n) \in \Pi^n$ and (likelihood) function $l : \mathcal{H} \rightarrow (0, 1]$ (Madansky (1964)). In particular, geometric opinion pools are externally Bayesian. An analogous concept can be defined for GOPs:

Definition 2 A GOP $G : \Pi^{2n} \rightarrow \Pi$ is called externally Bayesian if

$$G(p_1^l, \pi_1^l, \dots, p_n^l, \pi_n^l) = G(p_1, \pi_1, \dots, p_n, \pi_n)^l$$

for every profile $(p_1, \pi_1, \dots, p_n, \pi_n) \in \Pi^{2n}$ and (likelihood) function $l : \mathcal{H} \rightarrow (0, 1]$.

On the left hand side of this equation not only all posteriors are updated (π_i^l), but also all priors (p_i^l), because the incoming information is observed by everybody, hence part of the shared information, hence contained in the priors.

While external Bayesianity requires that it be irrelevant whether pooling happens before or after updating, a different question is whether it matters *who* in the group has observed a given information. *Internal Bayesianity* requires that it be irrelevant whether every or just a single person obtains a given information:

Definition 3 A GOP $G : \Pi^{2n} \rightarrow \Pi$ is called internally Bayesian if, for each person i ,

$$G(p_1, \pi_1, \dots, p_{i-1}, \pi_{i-1}, p_i, \pi_i^l, p_{i+1}, \pi_{i+1}, \dots, p_n, \pi_n) = G(p_1^l, \pi_1^l, \dots, p_n^l, \pi_n^l)$$

for every profile $(p_1, \pi_1, \dots, p_n, \pi_n) \in \Pi^{2n}$ and (likelihood) function $l : \mathcal{H} \rightarrow (0, 1]$.

On the left hand side of this equation, i 's prior is not updated (p_i , not p_i^l), because the incoming information, being observed just by person i , is not part of the shared information, hence not reflected in any prior. Internal Bayesianity is based on the idea that the collective probabilities should incorporate all information available *somewhere* in the group, whether it is held by a single or every person. External and internal Bayesianity together imply that, for each person i ,

$$G(p_1, \pi_1, \dots, p_{i-1}, \pi_{i-1}, p_i, \pi_i^l, p_{i+1}, \pi_{i+1}, \dots, p_n, \pi_n) = G(p_1, \pi_1, \dots, p_n, \pi_n)^l$$

for every profile $(p_1, \pi_1, \dots, p_n, \pi_n) \in \Pi^{2n}$ and (likelihood) function $l : \mathcal{H} \rightarrow (0, 1]$.

It turns out that, if a GOP G takes the form (9), then external and internal Bayesianity are in fact equivalent, and equivalent to external Bayesianity of F :

Theorem 2 If a generalised opinion pool $G : \Pi^{2n} \rightarrow \Pi$ has the form (9) where $F : \Pi^n \rightarrow \Pi$ is any opinion pool, the following conditions are equivalent:

- (i) G is externally Bayesian;
- (ii) G is internally Bayesian;
- (iii) F is externally Bayesian.

So, if one desires G to be externally or internally Bayesian, one is bound to use an externally Bayesian opinion pool F in our pooling formula (9), for instance a geometric opinion pool F , which leads to pooling formula (6), hence to (8) in the equal-weight case. There also exist more complex (non-geometric) externally Bayesian opinion pools F , characterised in full generality by Genest, McConway, and Schervish (1986, Theorem 2.5), but geometric ones become the only solutions if $|\mathcal{H}| \geq 3$ and F has some additional properties (see Genest, McConway, and Schervish (1986), Corollary 4.5).

Proof. I show that (i) is equivalent with each of (ii) and (iii). By (9),

$$G(p_1^l, \pi_1^l, \dots, p_n^l, \pi_n^l) \propto \frac{\pi_1^l}{p_1^l} \cdots \frac{\pi_n^l}{p_n^l} F(p_1^l, \dots, p_n^l),$$

and hence by (10)

$$G(p_1^l, \pi_1^l, \dots, p_n^l, \pi_n^l) \propto \frac{l\pi_1}{lp_1} \cdots \frac{l\pi_n}{lp_n} F(p_1^l, \dots, p_n^l) = \frac{\pi_1}{p_1} \cdots \frac{\pi_n}{p_n} F(p_1^l, \dots, p_n^l). \quad (11)$$

On the other hand, again by (9) and (10),

$$G(p_1, \pi_1, \dots, p_n, \pi_n)^l \propto l \frac{\pi_1}{p_1} \cdots \frac{\pi_n}{p_n} F(p_1, \dots, p_n) \propto \frac{\pi_1}{p_1} \cdots \frac{\pi_n}{p_n} F(p_1, \dots, p_n)^l. \quad (12)$$

Relations (11) and (12) together immediately imply that G is externally Bayesian if and only if F is externally Bayesian. Further, again by (9) and (10),

$$\begin{aligned} G(p_1, \pi_1, \dots, p_{i-1}, \pi_{i-1}, p_i, \pi_i^l, p_{i+1}, \pi_{i+1}, \dots, p_n, \pi_n) &\propto l \frac{\pi_1}{p_1} \dots \frac{\pi_n}{p_n} F(p_1, \dots, p_n) \\ &\propto \frac{\pi_1}{p_1} \dots \frac{\pi_n}{p_n} F(p_1, \dots, p_n)^l. \end{aligned}$$

This together with (11) implies that G is internally Bayesian if and only if F is externally Bayesian. ■

7 When is information independent, when not?

Let us go back to Theorem 1's assumption of Independent Information (Ind). This assumption is often a useful idealisation, even in situations where it fails. But what exactly are these real situations where (Ind) fails? An important source for failure is what I call *subgroup information*, that is, information held by more than one but less than all persons. I will prove that, under certain conditions, (Ind) holds *if and only if* there is no subgroup information.

By a person i 's *observation set* I mean, informally, the (possibly quite enormous) collection of i 's relevant observations/items of information. Formally, one may define i 's observation set as a set \mathcal{O}_i of non-empty *observations* $O \subseteq \Omega$.¹⁶ In the case of a jury faced with hypotheses about the defendant's guilt, i 's observation set might include the observations 'an insecure smile on the defendant's face', 'the defendant's fingerprint near the crime scene', 'two contradictory statements by witness x', etc.

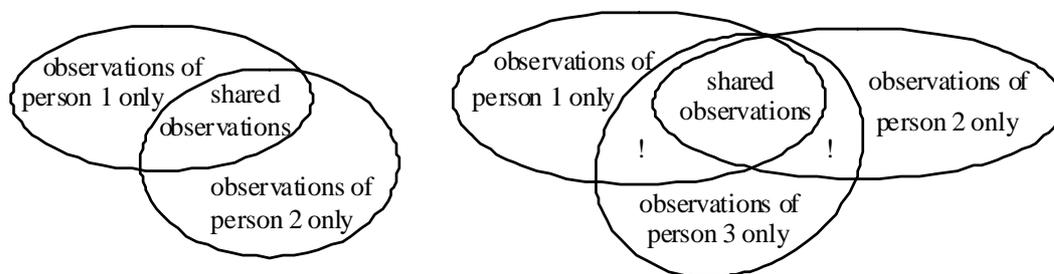


Figure 1: Observation sets in a group of $n = 2$ persons (no subgroup information), and a group of $n = 3$ persons (with subgroup information marked by "!!")

Figure 1 shows observation sets, *not* sets of possible worlds $A \subseteq \Omega$. These two concepts are in fact opposed to each other: the larger the observation set, the smaller the corresponding set of worlds (in which the observations hold); the *union* of

¹⁶An observation made by *every* person is represented by the sure event $O = \Omega$, because Ω is interpreted as containing the worlds that are possible under shared information. Formally, $O \in \mathcal{O}_1 \cap \dots \cap \mathcal{O}_n$ implies $O = \Omega$.

observation sets compares to the *intersection* of the sets of worlds. Formally, to an observation set \mathcal{O} corresponds the set of worlds $\cap_{O \in \mathcal{O}} O \subseteq \Omega$ (interpreted as Ω if $\mathcal{O} = \emptyset$). Thus i 's information E_i equals

$$E_i = \bigcap_{O \in \mathcal{O}_i \setminus (\mathcal{O}_1 \cap \dots \cap \mathcal{O}_n)} O,$$

the intersection of all of i 's observations except from any shared one; by footnote 16, this actually reduces to

$$E_i = \bigcap_{O \in \mathcal{O}_i} O.$$

Here is the problem. Consider any observation contained in the observation sets of more than one but less than all persons i – something impossible in groups of size $n = 2$ but possible in larger groups, as illustrated by the ‘!’ fields in Figure 1. This observation is not part of the shared information, but of the personal information E_i of *many* individuals i . Such subgroup information typically creates positive correlations between the E_i s in question. As a stylised example, consider a jury of $n = 3$ jurors faced with the hypothesis of guilt of the defendant (H). All jurors have read the charge (shared information), and moreover juror 1 has listened to the first witness report and observed the defendant’s nervousness (E_1), juror 2 has listened to the second witness report and observed the defendant’s smiles (E_2), and juror 3 has listened to both witness reports and had a private chat with the defendant (E_3). Note the subgroup information of jurors 1 and 3, and that of jurors 2 and 3, which typically causes E_3 to be positively correlated with E_1 and with E_2 . By contrast, individuals 1 and 2 together have no subgroup information. This situation is depicted in Figure 1 on the right.

To formally clarify the relationship between subgroup information and independence violation, some preparation is needed.

Definition 4 *A subgroup is a non-empty subset M of the group $N := \{1, \dots, n\}$. A subgroup is proper if it contains more than one but less than all persons.*

To formalise the notion of subgroup information, suppose that to each subgroup M there is a non-empty event $E^M \subseteq \Omega$, M 's *exclusively shared information*, representing all information held by *each of and only* the persons in M , where by assumption:

- $E_i = \bigcap_{\{i\} \subseteq M \subseteq N} E^M$ for all persons i (as i has observed those E^M with $i \in M$);¹⁷
- $E^N = \Omega$ (as *any* world $\omega \in \Omega$ is assumed possible under the shared information);

¹⁷Why not rather assume that $E_i = \bigcap_{\{i\} \subseteq M \subseteq N} E^M$, as E_i should not contain information held by *everybody*? In fact, both assumption are equivalent since by $E^N = \Omega$ an additional intersection with E^N has no effect.

- each E^M belongs to \mathcal{A} , the domain of the probability measure P (which holds in particular if \mathcal{A} contains *all* subsets of Ω).

For instance, the ‘!’ fields in Figure 1 (right) represent the observation sets corresponding to $E^{\{1,3\}}$ and $E^{\{2,3\}}$. E^M is interpretable as the intersection

$$\bigcap_{O \in (\cap_{i \in M} \mathcal{O}_i) \setminus (\cup_{i \notin M} \mathcal{O}_i)} O$$

of all observations O contained in *each* of the observation sets \mathcal{O}_i , $i \in M$, but in *none* of the observation sets \mathcal{O}_i , $i \notin M$ (where this intersection is Ω if $(\cap_{i \in M} \mathcal{O}_i) \setminus (\cup_{i \notin M} \mathcal{O}_i) = \emptyset$).

What we have to exclude is that a proper subgroup M exclusively shares information; in other words, E^M must be the no-information event Ω :

No Subgroup Information (NoSI) Every proper subgroup M has no exclusively shared information, i.e. $E^M = \Omega$ (or, more generally, $P(E^M) = 1$ ¹⁸).

This condition is empty if there are just $n = 2$ individuals, it requires $E^{\{1,2\}} = E^{\{1,3\}} = E^{\{2,3\}} = \Omega$ if $n = 3$, and it requires the ‘!’ fields in Figure 1 to be empty. Finally, consider the following independence assumption:

(Ind*) The events E^M , $\emptyset \neq M \subseteq N$, are (P -)independent conditional on each $H \in \mathcal{H}$.

(Ind*) is a more generally acceptable condition than (Ind) in that the E^M s, unlike the E_i s, are based on non-overlapping observation sets. Indeed, a subgroup M ’s exclusively shared information E^M , by the very meaning of ‘exclusively’, represents different observations than any other subgroup’s exclusively shared information.¹⁹

Theorem 3 *Assume (Ind*). Then:*

- Independent Information (Ind) is equivalent to No Subgroup Information (NoSI);*
- specifically, if $E^M \neq \Omega$ for proper subgroup M , then conditional on at least one $H \in \mathcal{H}$ the personal observations E_i , $i \in M$, are pairwise positively correlated (i.e. $P(E_i \cap E_j | H) > P(E_i | H)P(E_j | H)$ for any two distinct $i, j \in M$).*

¹⁸ $P(E^M) = 1$ is equivalent to ‘ $E^M = \Omega$ ’ in the natural case that only the empty event in \mathcal{A} has zero probability. Strictly speaking, ‘ $E^M = \Omega$ ’ means ‘no information’ while ‘ $P(E^M) = 1$ ’ means ‘essentially no information’. I am grateful to the referee for suggesting to require ‘ $P(E^M) = 1$ ’ instead of ‘ $E^M = \Omega$ ’, thereby making it possible to state Theorem 3 without assuming that only the empty event in \mathcal{A} has zero probability.

¹⁹ (Ind*) holds if the observations in $\mathcal{O}_1 \cup \dots \cup \mathcal{O}_n$ are mutually (conditionally) independent.

Proof. Suppose (Ind*). I prove part (a); the proof includes a proof of part (b).

(i) First, assume (NoSI). Each event $E^{\{i\}}$ coincides with E_i up to a set of probability zero, because

$$E_i = \bigcap_{\{i\} \subseteq M \subseteq N} E^M = E^{\{i\}} \cap \left(\bigcap_{\{i\} \subseteq M \subseteq N \text{ \& } |M| \geq 2} E^M \right),$$

in which $P\left(\bigcap_{\{i\} \subseteq M \subseteq N \text{ \& } |M| \geq 2} E^M\right) = 1$ by (NoSI). So, as the events $E^{\{1\}}, \dots, E^{\{n\}}$ are independent conditional on any $H \in \mathcal{H}$ by (Ind*), also the events E_1, \dots, E_n are independent conditional on any $H \in \mathcal{H}$.

(ii) Now assume (NoSI) is violated, and let M^* be a proper subgroup with $P(E^{M^*}) < 1$. I show that the events $E_i, i \in M^*$, are pairwise positively correlated conditional on at least one $H \in \mathcal{H}$ (which proves part (b) and also completes the proof of part (a) since E_1, \dots, E_n are then not independent conditional on that H). Consider any distinct $i, j \in M^*$. By $P(E^{M^*}) < 1$ there exists an $H \in \mathcal{H}$ with $P(E^{M^*}|H) < 1$. Since $E_i = \bigcap_{\{i\} \subseteq M \subseteq N} E^M$ and using (Ind*), we have

$$P(E_i|H) = \prod_{\{i\} \subseteq M \subseteq N} P(E^M|H).$$

The analogous argument for j yields

$$P(E_j|H) = \prod_{\{j\} \subseteq M \subseteq N} P(E^M|H).$$

So,

$$P(E_i|H)P(E_j|H) = \left[\prod_{\{i\} \subseteq M \subseteq N} P(E^M|H) \right] \times \left[\prod_{\{j\} \subseteq M \subseteq N} P(E^M|H) \right]. \quad (13)$$

Further, we have

$$\begin{aligned} E_i \cap E_j &= \left[\bigcap_{\{i\} \subseteq M \subseteq N} E^M \right] \cap \left[\bigcap_{\{j\} \subseteq M \subseteq N} E^M \right] \\ &= \left[\bigcap_{\{i\} \subseteq M \subseteq N} E^M \right] \cap \left[\bigcap_{\{j\} \subseteq M \subseteq N \setminus \{i\}} E^M \right]. \end{aligned}$$

So, by (Ind*),

$$P(E_i \cap E_j|H) = \left[\prod_{\{i\} \subseteq M \subseteq N} P(E^M|H) \right] \times \left[\prod_{\{j\} \subseteq M \subseteq N \setminus \{i\}} P(E^M|H) \right]. \quad (14)$$

The relations (13) and (14) together entail

$$P(E_i \cap E_j | H) > P(E_i | H)P(E_j | H),$$

because expression (13) equals expression (14) multiplied with the factor

$$\prod_{\{i,j\} \subseteq M \subseteq N} P(E^M | H),$$

which is smaller than 1 since it contains the term $P(E^{M^*} | H) < 1$. ■

8 Opinion pooling in the presence of subgroup information

One may always try to ‘remove’ subgroup information through active information sharing prior to aggregation: all proper subgroups with exclusively shared information communicate this information to the rest of the group. In Figure 1, the observations in each ‘!’ field are communicated to the third person, and in the above jury example the subgroups $\{1, 3\}$ and $\{2, 3\}$ communicate the exact content of the first resp. second witness report to the third juror. Having thus removed any subgroup information, (NoSI) and hence (in view of Theorem 3) Independent Information (Ind) hold, so that opinion pooling can proceed along the lines of Sections 2-5.

But suppose now that such information sharing is not feasible, e.g. due to the complexity of subgroup information. Then (NoSI) fails, and hence (Ind) fails, so that we need to modify our pooling formula. It is at first not obvious whether and how one can generalise Theorem 1 to arbitrary information overlaps, i.e. whether and how collective opinions can incorporate all information spread around the group. The generalisation is possible, as will be seen. Roughly speaking, we have to replace Theorem 1’s axioms of Individual Bayesian Rationality (IBR) and Independent Information (Ind) by corresponding axioms based on subgroups rather than individuals. Theorem 1’s two other axioms, (APLA) and (CP), will not anymore appear explicitly, but are build implicitly into the model, as explained in a moment. The adapted axioms will again lead to a unique collective opinion π , calculated in a somewhat more complicated way than in Theorem 1.

First, let me state the new model ingredients, and compare them to the earlier ingredients. As before, we have a non-empty set of possible worlds Ω , partitioned into a countable set \mathcal{H} of non-empty hypotheses H . While Theorem 1’s model contained for every individual i a personal information $E_i \subseteq \Omega$, now for every subgroup M there is a non-empty event $E^M \subseteq \Omega$, M ’s *exclusively shared information*, representing all information held by *each of and only* the persons in M . By assumption, $E^N = \Omega$, reflecting that *any* world $\omega \in \Omega$ is possible under the shared information. From these

events E^M we can define each individual i 's information as

$$E_i = \bigcap_{\{i\} \subseteq M \subseteq N} E^M,$$

representing all information held *at least* by person i .

The earlier model contained every individual i 's (prior) belief P_i ; this is not anymore needed here. Instead, I only assume a single probability measure P , defined on some σ -algebra $\mathcal{A} \subseteq \mathcal{P}(\Omega)$ containing each E^M and each hypothesis $H \in \mathcal{H}$. We interpret P as capturing *common* prior beliefs.²⁰ This assumption of common prior beliefs is a simplification; it for instance implies that conditions such as (CP) and (APLA) above are built into the model, and hence will not have to appear explicitly.

Recall further that in Theorem 1's model (in its common prior version) people provide individual opinions π_1, \dots, π_n (reflecting 'individually shared' information) and a common prior opinion p (reflecting the group's shared information). So, technically, the earlier model contained the opinions π_1, \dots, π_n, p reflecting the shared information of the *improper* subgroups $\{1\}, \dots, \{n\}, N$, respectively. Our new model adds to this the opinions reflecting the shared information of certain *proper* subgroups $M \subseteq N$. More precisely, in the new model at least those (proper or improper) subgroup which exclusively share information will need to provide an opinion. Formally, let \mathcal{M} be a set of subgroups, containing at least those (proper or improper) subgroups $M \subseteq N$ with exclusively shared information, i.e. with $E^M \neq \Omega$. Without loss of generality, let $N \in \mathcal{M}$.²¹ Each subgroup M in \mathcal{M} submits an opinion $\pi_M \in \Pi$, representing M 's probability assignments based on M 's shared information (shared information need not be exclusively shared, i.e. may be known to other persons too; see Definition 5 below). Theorem 1's model (in the common prior version) is the special case that $\mathcal{M} = \{\{1\}, \dots, \{n\}, N\}$ ($= \{M : M \text{ is an improper subgroup}\}$) with $\pi_{\{1\}} = \pi_1, \dots, \pi_{\{n\}} = \pi_n, \pi_N = p$. In the last section's jury example with $n = 3$ individuals, we may put

$$\mathcal{M} = \{\{1\}, \{2\}, \{3\}, \{1, 3\}, \{2, 3\}, \{1, 2, 3\}\}$$

because $\{1, 2\}$ has no exclusively shared information.

In practice, in addition to every individual i with $\{i\} \in \mathcal{M}$ submitting an opinion $\pi_{\{i\}}$, every non-singleton subgroup $M \in \mathcal{M}$ will have to 'sit together', find out about the information it shares, and come up with an opinion π_M based on this shared information.

The technique to calculate the (collective) opinion $\pi \in \Pi$ from the submitted subgroup opinions π_M , $M \in \mathcal{M}$, is recursive. Let me first illustrate it by an example.

²⁰More precisely, I do not mean to assume that every individual i holds a belief on *all* events in \mathcal{A} . Rather i holds beliefs (at least) on a sub- σ -algebra of \mathcal{A} containing all hypotheses in \mathcal{H} and those events E^M for which $i \in M$. i 's beliefs on this sub- σ -algebra are given by P .

²¹One may always define \mathcal{M} as containing *all* subgroups, but in practice this maximal choice adds unnecessary steps to the recursive pooling procedure introduced below. The minimal choice is $\mathcal{M} = \{M : \emptyset \neq M \subsetneq N \text{ and } E^M \neq \Omega\} \cup \{N\}$.

Example. As in the last section's jury example, let there be $n = 3$ individuals and let $\mathcal{M} = \{\{1\}, \{2\}, \{3\}, \{1, 3\}, \{2, 3\}, \{1, 2, 3\}\}$. So, functions $\pi_{\{1\}}, \pi_{\{2\}}, \pi_{\{3\}}, \pi_{\{1,3\}}, \pi_{\{2,3\}}$ and $\pi_{\{1,2,3\}}$ are submitted. The recursion works as follows, where I use a slightly simplified version of the later notation and give only informal justifications.

- First, merge $\pi_{\{1,3\}}$ and $\pi_{\{2,3\}}$ into a function $\pi_{\{1,3\},\{2,3\}}$ that combines $\{1, 3\}$'s shared information and $\{2, 3\}$'s shared information. One may apply Corollary 1's formula:

$$\pi_{\{1,3\},\{2,3\}} \propto \pi_{\{1,3\}}\pi_{\{2,3\}}/\pi_{\{1,2,3\}}.$$

(To see why $\pi_{\{1,2,3\}}$ can play the role of the prior opinion p in Corollary 1, recall that p there represents the information shared by all individual opinions. The information shared by the opinions $\pi_{\{1,3\}}$ and $\pi_{\{2,3\}}$ is the information held by [1 and 3] *and* by [2 and 3]. This is equivalent to the information held by 1 and 2 and 3, i.e. the information expressed in $\pi_{\{1,2,3\}}$.)

- Next, define $\pi_{\{1,2\}}$ as $\pi_{\{1,2,3\}}$, because the subgroup $\{1, 2\}$ does not exclusively share any information and hence shares the same information as the larger group $\{1, 2, 3\}$.
- Next, merge $\pi_{\{1\}}$ and $\pi_{\{2\}}$ into a function $\pi_{\{1\},\{2\}}$ that combines $\{1\}$'s and $\{2\}$'s information. One may apply Corollary 1's formula:

$$\pi_{\{1\},\{2\}} \propto \pi_{\{1\}}\pi_{\{2\}}/\pi_{\{1,2\}}.$$

(Why can $\pi_{\{1,2\}}$ play the role of p in Corollary 1, i.e. why does $\pi_{\{1,2\}}$ express the information shared by $\pi_{\{1\}}$ and $\pi_{\{2\}}$? The information shared by $\pi_{\{1\}}$ and $\pi_{\{2\}}$ is the information held by 1 and by 2, i.e. the information expressed in $\pi_{\{1,2\}}$.)

- Finally, merge $\pi_{\{1\},\{2\}}$ and $\pi_{\{3\}}$ into the function $\pi = \pi_{\{1\},\{2\},\{3\}}$ that combines $\{1\}$'s, $\{2\}$'s and $\{3\}$'s information. Again, one may apply Corollary 1's formula:

$$\pi = \pi_{\{1\},\{2\},\{3\}} \propto \pi_{\{1\},\{2\}}\pi_{\{3\}}/\pi_{\{1,3\},\{2,3\}}.$$

(Why can $\pi_{\{1,3\},\{2,3\}}$ play the role of p in Corollary 1, i.e. why does $\pi_{\{1,3\},\{2,3\}}$ represent the information shared by $\pi_{\{1\},\{2\}}$ and $\pi_{\{3\}}$? The information shared by $\pi_{\{1\},\{2\}}$ and $\pi_{\{3\}}$ is the information held by [1 or 2] and by 3. This is precisely the information held by [1 and 3] or by [2 and 3], i.e. the information expressed in $\pi_{\{1,3\},\{2,3\}}$.)

Now I come to the formal treatment. Recall that i 's information E_i is given by

$$E_i = \bigcap_{\{i\} \subseteq M \subseteq N} E^M,$$

i.e. i knows precisely the conjunction of what the subgroups containing i exclusively share. This generalises as follows to:

Definition 5 A subgroup M 's shared information is defined as

$$E_M := \bigcap_{M \subseteq M' \subseteq N} E^{M'}$$

(the conjunction of all information exclusively shared by some supergroup of M).

E_M represents what is known to *at least* all members of M – as opposed to M 's *exclusively* shared information E^M , known *exactly* all members of M . Taking the case of a singleton subgroup $M = \{i\}$, the event $E_{\{i\}}$ coincides with E_i . Also, note that

$$P(E^M) > 0 \text{ and } P(E_M) > 0 \text{ for each subgroup } M$$

because

$$P(E^M), P(E_M) \geq P\left(\bigcap_{\emptyset \neq M' \subseteq N} E^{M'}\right) = P(E_1 \cap \dots \cap E_n) > 0.$$

The following condition translates Individual Bayesian Rationality (IBR) to subgroups in \mathcal{M} :

Subgroup Bayesian Rationality (SBR) $\pi_M(H) = P(H|E_M)$ for every subgroup $M \in \mathcal{M}$ and hypothesis $H \in \mathcal{H}$.

As in Theorem 1, we would like the collective opinion to satisfy Collective Bayesian Rationality (CBR); that is, we require that

$$\pi(H) = P(H|E_1 \cap \dots \cap E_n) \text{ for each hypothesis } H \in \mathcal{H},$$

a condition that may be rewritten in several equivalent ways since (by Definition 5)

$$E_1 \cap \dots \cap E_n = E_{\{1\}} \cap \dots \cap E_{\{n\}} = \bigcap_{\emptyset \neq M \subseteq N} E^M = \bigcap_{\emptyset \neq M \subseteq N} E_M.$$

As a technical tool to construct collective opinion π satisfying (CBR), I need to introduce opinions of *abstract* individuals.

Definition 6 An abstract individual is a non-empty set A of subgroups M ; its order is $order(A) := \min\{|M| : M \in A\}$, the size of a smallest subgroup in A .

The opinions $\pi_{\{1,3\},\{2,3\}}, \pi_{\{1\},\{2\}}, \dots$ defined in the example above are in fact the opinions of the abstract individuals $\{\{1, 3\}, \{2, 3\}\}, \{\{1\}, \{2\}\}, \dots$. More generally, I interpret an abstract individual A as a hypothetical agent who knows the shared information of any subgroup $M \in A$ (and no more). For instance, $A = \{\{1, 3\}, \{2, 3\}\}$ knows $\{1, 3\}$'s shared information *and* $\{2, 3\}$'s shared information. A 's information is thus given by $\bigcap_{M \in A} E_M$. To get a concrete idea, note that the abstract agent $A = \{\{1, 3\}, \{2, 3\}\}$ knows

- at least as much as the abstract agent $\{\{1, 3\}\}$, who knows all information that 1 and 3 share (but no information that 2 and 3 share exclusively);
- at least as much as the abstract agent $\{\{1, 2, 3\}\}$, who knows all information that 1, 2 and 3 share (but no information that 1 and 3 share exclusively or that 2 and 3 share exclusively);
- at most as much as the abstract agent $\{\{1\}, \{2\}, \{3\}\}$, who knows all that 1 or 2 or 3 knows (and hence all that two or three of these individuals know together).

I will calculate for each abstract individual A an opinion $\pi_A \in \Pi$ that reflects precisely A 's information $\bigcap_{M \in A} E_M$, i.e. that satisfies

$$\pi_A(H) = P\left(H \mid \bigcap_{M \in A} E_M\right) \text{ for each } H \in \mathcal{H}. \quad (15)$$

Specifically, I calculate π_A by backward recursion over $order(A)$: π_A is calculated first for $order(A) = n$, then for $order(A) = n - 1, \dots$, then for $order(A) = 1$. This finally yields π , since by (CBR) and (15)

$$\pi = P(\cdot \mid E_{\{1\}} \cap \dots \cap E_{\{n\}}) = \pi_A$$

where A is the abstract individual $\{\{1\}, \{2\}, \dots, \{n\}\}$ of order 1. In the recursive construction, the main steps are to calculate from opinions π_A and π_{A^*} of abstract individuals A and A^* the opinion $\pi_{A \cup A^*}$ of the abstract individual $A \cup A^*$ whose information combines the information of A and A^* . To derive $\pi_{A \cup A^*}$ from π_A and π_{A^*} , I generalise the formula of Theorem 1 to (two) *abstract* individuals. To do so, the notion of shared information is crucial. What information do A and A^* share? They share precisely the information held by the abstract individual

$$A \vee A^* := \{M \cup M^* : M \in A \text{ and } M^* \in A^*\}.$$

The reason is: the information A and A^* share is precisely the information that A knows *and* A^* knows, i.e. that some subgroup in A shares *and* some subgroup in A^* shares, i.e. that some union $M \cup M^*$ with $M \in A$ and $M^* \in A^*$ shares. So, when combining opinions π_A and π_{A^*} , $A \vee A^*$'s opinion $\pi_{A \vee A^*}$ plays the role of the common prior p in Theorem 1. More precisely, the crucial result on how to combine opinions of abstract individuals states as follows (and is proved later):

Lemma 1 *Assume (Ind^{*}). Consider abstract individuals B and C , form the abstract individuals $B \vee C$ and $B \cup C$. If $\pi_B, \pi_C, \pi_{B \vee C}$ are opinions in Π given by (15), then*

- *there is an opinion in Π proportional to $\pi_B \pi_C / \pi_{B \vee C}$,*²²

²²Equivalently, the sum $\sum_{H \in \mathcal{H}} \pi_B(H) \pi_C(H) / \pi_{B \vee C}(H)$ is finite. Indeed, a function f from \mathcal{H} to $(0, \infty)$ (such as $\pi_B \pi_C / \pi_{B \vee C}$) can be normalised to a function with sum 1 if and only if f has a finite sum.

- this opinion is the function $\pi_{B \cup C}$ given by (15).

The formula in Lemma 1 guides us in assigning opinions to abstract individuals. The assignment is recursive, with another nested recursion in ‘Case 2’:

Definition 7 Define the opinions $\pi_A \in \Pi$ of abstract individual A by the following backward recursion on $\text{order}(A)$:

- Assume $\text{order}(A) = n$. Then $A = \{N\}$. Define $\pi_A := \pi_N$.
- Assume $\text{order}(A) = k < n$ and assume $\pi_{A'}$ is already defined for $\text{order}(A') > k$.

Case 1: $|A| = 1$. Then $A = \{M\}$. If $M \in \mathcal{M}$, define $\pi_A = \pi_M$. If $M \notin \mathcal{M}$, consider the abstract individual $A' := \{M \cup \{i\} : i \notin M\}$ containing all subgroups with exactly one person added to M (interpretation: A and A' have the same information by $M \notin \mathcal{M}$) and define $\pi_A := \pi_{A'}$ (where $\pi_{A'}$ is already defined by $\text{order}(A') = k + 1$).

Case 2: $|A| > 1$. Define π_A by another recursion on $|\{M \in A : |M| = k\}|$, the number of subgroups in A of size k :

- Assume $|\{M \in A : |M| = k\}| = 1$. Then $A = \{M\} \cup A^*$, where $|M| = k$ and $\text{order}(A^*) > k$. Define π_A by $\pi_A \propto \pi_{\{M\}} \pi_{A^*} / \pi_{\{M\} \vee A^*}$ (where $\pi_{\{M\}}$ is already defined in case 1, and π_{A^*} and $\pi_{\{M\} \vee A^*}$ are already defined by $\text{order}(A^*) > k$ and $\text{order}(\{M\} \vee A^*) > k$).
- Assume $|\{M \in A : |M| = k\}| = l > 1$ and assume π_{A^*} is already defined for all the A^* such that $|\{M \in A^* : |M| = k\}| < l$ (and $\text{order}(A^*) = k$). Then $A = \{M\} \cup A^*$ with $|M| = k$ and $|\{M^* \in A^* : |M^*| = k\}| = l - 1$. Define π_A by $\pi_A \propto \pi_{\{M\}} \pi_{A^*} / \pi_{\{M\} \vee A^*}$ (where $\pi_{\{M\}}$ is already defined in case 1, π_{A^*} is already defined by $|\{M^* \in A^* : |M^*| = k\}| = l - 1$, and $\pi_{\{M\} \vee A^*}$ is already defined by $\text{order}(\{M\} \vee A^*) > k$).

On the last recursion step we reach the opinions π_A of abstract individuals of order 1, hence in particular the opinion of $A = \{\{1\}, \dots, \{n\}\}$, and this is the desired opinion that incorporates the group’s full information:

Theorem 4 If subgroups satisfy (SBR), information satisfies (Ind*), and the collective satisfies (CBR), then the collective opinion π is given by $\pi_{\{\{1\}, \dots, \{n\}\}}$, the (recursively calculated) opinion of the abstract individual $\{\{1\}, \dots, \{n\}\}$.

The procedure needed to obtain the collective opinion π ($= \pi_{\{\{1\}, \dots, \{n\}\}}$) may have a high complexity.²³ How practically feasible is it? One should distinguish two separate tasks: (i) first, each subgroup M in \mathcal{M} has to form and submit an opinion

²³I am grateful to the referee for drawing my attention to this point.

π_M ; (ii) subsequently, the collective opinion π has to be derived algorithmically from the various subgroup opinions. Let me comment on each task.

Whether task (i) is feasible in practice depends crucially on the number and size of subgroups in \mathcal{M} , which in turn depends on how information is distributed across people. In the worst case, *every* subgroup exclusively shares information. Here, $\mathcal{M} = \mathcal{P}(N) \setminus \{\emptyset\}$ and $|\mathcal{M}| = 2^n - 1$, and the task becomes infeasible already for moderately large n . On the other hand, the task seems more feasible in situations where only relatively few subgroups exclusively share information. Suppose for instance that, when pooling expert opinions relative to certain hypotheses about climate change, only the following subgroups exclusively share information: each single expert, i.e. each singleton subgroup $\{i\} \subseteq N$; a group of physicists $M_1 \subseteq N$; a group of biologists $M_2 \subseteq N$; and a group of meteorologists $M_3 \subseteq N$. Then we may define \mathcal{M} as $\{\{1\}, \{2\}, \dots, \{n\}, M_1, M_2, M_3, N\}$, so that only $|\mathcal{M}| = n + 4$ opinions have to be formed and submitted.

Task (ii) involves an algorithm with a nested recursion; the overall number of steps grows more than exponentially in n .²⁴ So, for large n , task (ii) poses a feasibility problem – even if $|\mathcal{M}|$ is small, i.e., if task (i) seems feasible. There is however an escape to this problem if \mathcal{M} contains only relatively small proper subgroups. Let $m := \max_{M \in \mathcal{M} \setminus \{N\}} |M|$ denote the maximal size that subgroups in \mathcal{M} can have (apart from the improper subgroup N). A quick inspection of the algorithm in Definition 7 shows that its backward recursion (which assigns opinions π_A to abstract agents A) is trivial until it reaches abstract agents of order m : all abstract agents of order $k > m$ get assigned the opinion $\pi_A = \pi_N$. So a shortcut is possible: define π_A as π_N for all abstract agents A with $\text{order}(A) > m$, and start the backward recursion with those abstract agents A with $\text{order}(A) = m$.

Even if both tasks (i) and (ii) turn out to be practically feasible, the very choice of \mathcal{M} (before starting task (i)) may pose another high-complexity problem. Suppose \mathcal{M} is chosen by surveying all subgroups one by one to find out which ones exclusively share information (each subgroup might be asked to ‘sit together’ and look for potential information overlaps). Since there are $2^n - 1$ subgroups in total, this would become infeasible already for moderately large n . However, no such problem arises if \mathcal{M} can be specified without performing an explicit subgroup-by-subgroup examination. For instance, \mathcal{M} might be specified by a social planner who knows from the start that certain subgroups (say, those containing experts from different fields) do not exclusively share any information, while the other subgroups might exclusively share information.

Turning now to the proof, I first show Lemma 1 and then Theorem 4.

Proof of Lemma 1. Assume (Ind*). Let B, C be abstract individuals, and π_B, π_C ,

²⁴In the algorithm, for each abstract agent A an opinion π_A is calculated. There are $2^{2^n - 1} - 1$ abstract agents in total. Hence, $2^{2^n - 1} - 1$ opinions have to be calculated.

$\pi_{B \vee C}, \pi_{B \cup C} \in \Pi$. Suppose $\pi_B, \pi_C, \pi_{B \vee C}$ satisfy (15). For all abstract individuals A , put

$$\bar{A} := \{M \subseteq N : M' \subseteq M \text{ for some } M' \in A\},$$

the set of supergroups of subgroups in A . By (15), $\pi_{B \vee C} = P(\cdot | \bigcap_{M \in B \vee C} E_M)$, where by Definition 5

$$\bigcap_{M \in B \vee C} E_M = \bigcap_{M \in B \vee C} \bigcap_{M \subseteq M' \subseteq N} E^{M'} = \bigcap_{M \in \overline{B \vee C}} E^M.$$

So,

$$\pi_{B \vee C} = P(\cdot | E) \text{ with } E := \bigcap_{M \in \overline{B \vee C}} E^M. \quad (16)$$

Analogously, by (15), $\pi_B = P(\cdot | \bigcap_{M \in B} E_M)$, where by Definition 5

$$\bigcap_{M \in B} E_M = \bigcap_{M \in B} \bigcap_{M \subseteq M' \subseteq N} E^{M'} = \bigcap_{M \in \bar{B}} E^M = E_B \cap E$$

with $E_B := \bigcap_{M \in \bar{B} \setminus \overline{B \vee C}} E^M$. So $\pi_B = P(\cdot | E_B \cap E)$, and hence by Bayes' rule

$$\pi_B \propto P(\cdot | E) P(E_B | \cdot \cap E). \quad (17)$$

By an analogous argument for C , we have

$$\pi_C \propto P(\cdot | E) P(E_C | \cdot \cap E), \quad (18)$$

where $E_C := \bigcap_{M \in \bar{C} \setminus \overline{B \vee C}} E^M$. By (16), (17) and (18) we have

$$\begin{aligned} \pi_B \pi_C / \pi_{B \vee C} &\propto [P(\cdot | E) P(E_B | \cdot \cap E)] [P(\cdot | E) P(E_C | \cdot \cap E)] / P(\cdot | E) \\ &= P(\cdot | E) P(E_B | \cdot \cap E) P(E_C | \cdot \cap E). \end{aligned} \quad (19)$$

Note that each of E_B, E_C, E is an intersection of a set of events of type E^M , where the three sets of E^M s (corresponding to E_B, E_C, E , respectively) are pairwise disjoint. So, as by (Ind*) all E^M s are independent conditional on any $H \in \mathcal{H}$, so are the events E_B, E_C, E . Consider an $H \in \mathcal{H}$. As E_B, E_C, E are independent given H , the events E_B, E_C are independent given $H \cap E$. So

$$P(E_B | \cdot \cap E) P(E_C | \cdot \cap E) = P(E_B \cap E_C | \cdot \cap E).$$

Substituting this into (19) and then applying Bayes' rule, we obtain

$$\pi_B \pi_C / \pi_{B \vee C} \propto P(\cdot | E) P(E_B \cap E_C | \cdot \cap E) \propto P(\cdot | E_B \cap E_C \cap E) \in \Pi.$$

Now suppose $\pi_{B \cup C} = P(\cdot | E_B \cap E_C \cap E)$. We may rewrite $E_B \cap E_C \cap E$ as

$$\bigcap_{M \in \overline{B \cup C}} E^M = \bigcap_{M \in B \cup C} \bigcap_{M \subseteq M' \subseteq N} E^{M'} = \bigcap_{M \in B \cup C} E^M,$$

and hence π_{BUC} equals $P(\cdot | \bigcap_{M \in BUC} E_M)$, i.e. satisfies (15). ■

Proof of Theorem 4. Assume (SBR) and (Ind*). By backward induction on the order of A I show that each abstract individual A has opinion π_A satisfying (15). This in particular implies that $\{\{1\}, \dots, \{n\}\}$ has opinion

$$\pi_{\{\{1\}, \dots, \{n\}\}}(H) = P(H | E_1 \cap \dots \cap E_n) \text{ for each } H \in \mathcal{H},$$

so that under (CBR) we have $\pi = \pi_{\{\{1\}, \dots, \{n\}\}}$, as desired.

Denote by \mathbf{A} the set of abstract individuals A . The recursion proceeds as follows.

- If $order(A) = n$, then $A = \{N\}$, and by definition $\pi_A = \pi_N$. So by (SBR) $\pi_A = P(\cdot | E_N) = P(\cdot | \bigcap_{M \in A} E_M)$, as desired.
- Now let $order(A) = k < n$, and assume (15) holds for all $A' \in \mathbf{A}$ with $order(A') > k$. I have to show that $\pi_A = P(\cdot | \bigcap_{M \in A} E_M)$.

Case 1: $|A| = 1$. Then $A = \{M\}$ with $|M| = k$. If $M \in \mathcal{M}$, then by definition $\pi_A = \pi_M$, so by (SBR) $\pi_A = P(\cdot | E_M) = P(\cdot | \bigcap_{M' \in A} E_{M'})$, as desired. Now assume $M \notin \mathcal{M}$. Then by definition $\pi_A = \pi_{A'}$ with $A' := \{M \cup \{i\} : i \notin M\}$. Since $order(A') = k + 1$, the induction hypothesis yields $\pi_{A'} = P(\cdot | \bigcap_{M' \in A'} E_{M'})$, hence $\pi_A = P(\cdot | \bigcap_{M' \in A'} E_{M'})$. So I have to show that $\bigcap_{M' \in A'} E_{M'} = E_M$. By Definition 5,

$$E_M = \bigcap_{M \subseteq M' \subseteq N} E_{M'} = E^M \cap \left\{ \bigcap_{M' \in A'} \left[\bigcap_{M' \subseteq M'' \subseteq N} E_{M''} \right] \right\}.$$

In this, $E^M = \Omega$ (by $M \notin \mathcal{M}$) and $\bigcap_{M' \subseteq M'' \subseteq N} E_{M''} = E_{M'}$ (by Definition 5). So $E_M = \bigcap_{M' \in A'} E_{M'}$, as desired.

Case 2: $|A| > 1$. I show $\pi_A = P(\cdot | \bigcap_{M \in A} E_M)$ by induction on the number $|\{M \in A : |M| = k\}|$ of subgroups in A of size k .

- Let $|\{M \in A : |M| = k\}| = 1$. Then $A = \{M\} \cup A^*$ with $|M| = k$ and $order(A^*) > k$. Then π_A was defined as the function in Π proportional to $\pi_{\{M\}} \pi_{A^*} / \pi_{\{M\} \vee A^*}$; let me show that (i) such a function does indeed exist and (ii) satisfies (15), as desired. Now, $\pi_{\{M\}}$ satisfies (15) by Case 1, and π_{A^*} and $\pi_{\{M\} \vee A^*}$ satisfy (15) by $order(A^*) > k$ and $order(\{M\} \vee A^*) > k$ (and the k -induction hypothesis). So, by Lemma 1, the function $\pi_{\{M\}} \pi_{A^*} / \pi_{\{M\} \vee A^*}$ is proportional to a function in Π , so that π_A is well-defined. Also by Lemma 1, this function π_A satisfies (15), as desired.
- Let $|\{M \in A : |M| = k\}| = l > 1$, and assume A^* satisfies (15) whenever $|\{M \in A^* : |M| = k\}| < l$ (and $order(A^*) = k$). By definition, $\pi_A \propto \pi_{\{M\}} \pi_{A^*} / \pi_{\{M\} \vee A^*}$, where $A = \{M\} \cup A^*$ with $|M| = k$ and $|\{M^* \in A^* :$

$|M^*| = k\} = l - 1$. Again, we have to show that π_A is well-defined (i.e. that Π indeed contains a function proportional to $\pi_{\{M\}}\pi_{A^*}/\pi_{\{M\}\vee A^*}$) and satisfies (15). $\pi_{\{M\}}$ satisfies (15) by Case 1, π_{A^*} satisfies (15) by $|\{M^* \in A^* : |M^*| = k\}| = l - 1$ (and the l -induction hypothesis), and $\pi_{\{M\}\vee A^*}$ satisfies (15) by $order(\{M\} \vee A^*) > k$ (and the k -induction hypothesis). So, by Lemma 1, π_A is well-defined and satisfies (15). ■

9 Conclusion

The above model interprets opinion pooling as information pooling: collective opinions should build in the group's entire information, be it shared or personal. According to the pooling formulae I obtained, collective opinions should account for informational asymmetries *not* by taking a standard weighted (linear or geometric) average of the individual opinions with higher weight assigned to better informed individuals *but* by incorporating people's prior opinions in addition to their actual (i.e. posterior) opinions. In practice, people have either to agree on a common prior opinion p , i.e. on how to interpret the shared information, or they have to submit their possibly diverging prior opinions p_1, \dots, p_n . Based on simple axioms, Theorem 1 shows how to aggregate the (prior and posterior) opinions into a collective opinion. The formula defines a *multiplicative* opinion pool: the collective opinion π is the product of the individual opinions π_1, \dots, π_n and a function g (which depends on prior opinions).

More precisely, Theorem 1 suggests that, based on individual opinions π_1, \dots, π_n , the collective opinion π should be defined by

$$\pi \propto \pi_1 \cdots \pi_n / p^{n-1}$$

if people agree on a common prior p , and by

$$\pi \propto \frac{\pi_1}{p_1} \cdots \frac{\pi_n}{p_n} F(p_1, \dots, p_n) \quad (20)$$

if people have arbitrary priors p_1, \dots, p_n , where F is a standard opinion pool. I have suggested that F should be anonymous (i.e. symmetric in its arguments) because the prior opinions it pools are based on the same (shared) information, giving no individual an informational superiority. More specifically, I have suggested to define F as unweighted geometric pooling, because this generates appealing properties shown in Theorem 2. This choice of F gives collective opinion the form

$$\pi \propto \frac{\pi_1}{p_1^{1-1/n}} \cdots \frac{\pi_n}{p_n^{1-1/n}}.$$

Fortunately, not much depends on how we choose F in the pooling formula (20) if – as is frequently the case – $\frac{\pi_1}{p_1} \cdots \frac{\pi_n}{p_n}$ dominates $F(p_1, \dots, p_n)$ (i.e., if the function $\frac{\pi_1}{p_1} \cdots \frac{\pi_n}{p_n}$ varies far more than the function $F(p_1, \dots, p_n)$ for ‘reasonable’ choices of F).

In such cases, one might in practice refrain from choosing F and simply define the collective opinion as

$$\pi \propto \frac{\pi_1}{p_1} \dots \frac{\pi_n}{p_n},$$

a particularly elegant pooling formula.

A crucial axiom underlying these pooling formulas is that personal information is independent. By Theorem 3, independence is threatened by the possibility of subgroup information, i.e. of information held by more than one but less than all individuals. Theorem 4 therefore generalises the aggregation rule to arbitrary information distributions (allowing for subgroup information). The generalisation is unique, but assumes that each subgroup with subgroup information agrees on how to interpret this information, a kind of common prior assumption. Dropping this assumption would have gone beyond the scope of this paper, but it might be an interesting route for future research.

10 References

- Bacharach, M. (1972) Scientific disagreement, *unpublished manuscript*
- Dietrich, F. (2004) Opinion Pooling under Asymmetric Information, working paper, *Public Economics 0407002, EconWPA*
- Dietrich, F. (2008) The premises of Condorcet's jury theorem are not simultaneously justified, *Episteme - a Journal of Social Epistemology* 5(1): 56-73
- Dietrich, F., and C. List (2007) Opinion pooling on general agendas, working paper, *METEOR Research Memorandum 038, Maastricht University*
- Fitelson, B. (2001) A Bayesian account of independent evidence with applications, *Philosophy of Science* 68 (Proceedings). S123 - S140
- Genest, C. (1984) A characterization theorem for externally Bayesian groups, *Ann. Statist.* 12, p. 1100-1105
- Genest, C. and J. V. Zidek (1986) Combining probability distributions: a critique and an annotated bibliography, *Statist. Sci.* 1, p. 114-148
- Genest, C., K. J. McConway and M. J. Schervish (1986) Characterization of externally Bayesian pooling operators, *Ann. Statist.* 14, 487-501
- Hild, M. (1998) The instability of ex post aggregation, *Typescript*
- Hylland, A. and R. Zeckhauser (1979) The impossibility of group decision making with separate aggregation of beliefs and values, *Econometrica* 47, p. 1321-36

- Jeffrey, R. (1983) (first published 1965) *The logic of decision*, Chicago: Chicago University Press
- Lehrer, K. and C. Wagner (1981) *Rational Consensus in Science and Society*, Dordrecht: Reidel
- Levi, I. (1990) Pareto-unanimity and consensus, *Journal of Philosophy* 87
- Madansky, A. (1964) Externally Bayesian groups, *Technical Report* RM-4141-PR, RAND Corporation
- McConway, K. (1978) The combination of experts' opinions in probability assessments: some theoretical considerations, *Ph.D. thesis*, University College London
- McConway, K. (1981) Marginalization and linear opinion pools, *Jour. Amer. Statist. Assoc.* 76, p. 410-414
- Mongin, P. (1995) Consistent Bayesian aggregation, *Journal of Economic Theory* 66, p. 313-351
- Mongin, P. (1998) The paradox of the Bayesian experts and state-dependent utility theory, *Journal of Mathematical Economics* 29, p. 331-61
- Morris, P. A. (1974) Decision analysis expert use, *Management Science* 20, p. 1233-41
- Pearl, J. (2000) *Causality: Models, Reasoning and Inference*, Cambridge: Cambridge University Press
- Pivato, M. (2008) The Discursive Dilemma and Probabilistic Judgement Aggregation, *MPRA Paper 8412*, University Library of Munich, Germany
- Risse, M. (2001) Instability of ex post aggregation in the Bolker/Jeffrey framework and related instability phenomena, *Erkenntnis* 55, p. 239-269
- Risse, M. (2003) Bayesian group agents and two modes of aggregation, *Synthese*, forthcoming
- Savage, L. (1954) *The foundations of statistics*, New York: Wiley
- Schervish, M., T. Seidenfeld and J. Kadane (1991) Shared preferences and state-dependent utilities, *Management Science* 37, p. 1575-89
- Seidenfeld, T., J. Kadane and M. Schervish (1989) On the shared preferences of two Bayesian decision makers, *Journal of Philosophy* 86, p. 221-44
- Wagner, C. G. (1982) Allocation, Lehrer models, and the consensus of probabilities, *Theory and Decision* 14, p. 207-220

Wagner, C. (1985) "On the Formal Properties of Weighted Averaging as a Method of Aggregation", *Synthese* 62: 97-108